

IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF TEXAS
WACO DIVISION

VERVAIN, LLC,

Plaintiff,

v.

MICRON TECHNOLOGY, INC., MICRON
SEMICONDUCTOR PRODUCTS, INC.,
AND MICRON TECHNOLOGY TEXAS,
LLC,

Defendants.

Case No. 6:21-cv-487-ADA

JURY TRIAL DEMANDED

VERVAIN, LLC,

Plaintiff,

v.

WESTERN DIGITAL CORPORATION;
WESTERN DIGITAL TECHNOLOGIES,
INC.; AND HGST, INC.

Defendants.

Case No. 6:21-cv-488-ADA

JURY TRIAL DEMANDED

DEFENDANTS' OPENING CLAIM CONSTRUCTION BRIEF

TABLE OF CONTENTS

I.	Introduction.....	1
II.	Background.....	1
III.	Argument	4
A.	“SLC non-volatile memory” (298 patent, cls. 1, 5; 385 patent cls. 1, 5; 240 patent, cls. 1, 6); “SLC nonvolatile memory” (300 patent, cls. 1, 3, 4, 7, 12)	4
B.	“MLC non-volatile memory” (298 patent, cls. 1, 4; 385 patent, cls. 1, 4; 240 patent, cls. 1, 6); “MLC nonvolatile memory” (300 patent, cls. 1, 3, 4, 12)	4
C.	“data integrity test” (298 patent, cl. 1; 385 patent, cl. 1; 240 patent, cls. 1, 6; 300 patent, cls. 1, 12).....	8
D.	“comparing the stored data to the retained data in the random access volatile memory” (300 patent, cls. 1 and 12).....	11
E.	“to achieve enhanced endurance” (300 patent, cls. 1, 12)	15
F.	“the list of logical address ranges having a minimum quanta of addresses” (298 patent, cl. 1; 385 patent, cl. 1; 240 patent, cls. 1, 6)	19
G.	“wherein the controller is further adapted to determine which of the blocks of the plurality of the blocks in the MLC and SLC non-volatile memory modules are accessed most frequently and wherein the controller segregates those blocks that receive frequent writes into the at least one SLC non-volatile memory module and those blocks that receive infrequent writes into the at least one MLC nonvolatile module, and maintain a count value of the blocks in the MLC non-volatile memory module determined to have received frequent writes and that are accessed most frequently on a periodic basis when the count value is a predetermined count value, transfer the contents of the counted blocks in the MLC non-volatile memory module determined to have received frequent writes after reaching the predetermined count value to the SLC non-volatile memory module and which determined blocks in the SLC are determined in accordance with the next equivalent range of physical addresses determined by the controller” (240 patent, cl. 1)	21
H.	“wherein the controller is further adapted to maintain a count value of those blocks that are accessed most frequently and, on a periodic basis when the count value is a predetermined count value, transfer the contents of those counted blocks into the SLC non-volatile memory module, wherein the counted blocks transferred to after reaching the predetermined	

count value are determined in accordance with the next equivalent range of physical addresses determined by the controller” (240 patent, cl. 6)	25
I. “wherein the mapping is performed as necessitated by the system to maximize lifetime” (300 patent, cls. 1, 12).....	26
J. “wherein a failure of the data integrity test performed by the controller results in a remapping of the address space to a different physical range of addresses and transfer of data corresponding to the stored data to those remapped physical addresses” (300 patent, cls. 1, 12)	28
IV. Conclusion	30

TABLE OF AUTHORITIES

Cases

<i>Ajinomoto Co. v. ITC</i> , 932 F.3d 1342 (Fed. Cir. 2019).....	14
<i>Alcon Rsch., Ltd. v. Apotex Inc.</i> , 687 F.3d 1362 (Fed. Cir. 2012).....	19
<i>Becton, Dickinson & Co. v. Tyco Healthcare Grp., LP</i> , 616 F.3d 1249 (Fed. Cir. 2010).....	5
<i>Bell Atl. Network Servs., Inc. v. Covad Commc'n Group, Inc.</i> , 262 F.3d 1258 (Fed. Cir. 2001).....	5
<i>Berkheimer v. HP Inc.</i> , 881 F.3d 1360 (Fed. Cir. 2018).....	19, 26
<i>Biogen Idec, Inc. v. GlaxoSmithKline LLC</i> , 713 F.3d 1090 (Fed. Cir. 2013).....	14
<i>CAE Screenplates, Inc. V. Henrich Fiedler GMBH & Co. KG</i> , 224 F.3d 1308 (Fed. Cir. 2000).....	5
<i>Chef Am., Inc. v. Lamb-Weston, Inc.</i> , 358 F.3d 1371 (Fed. Cir. 2004).....	30
<i>Engel Indus., Inc. v. Lockformer Co.</i> , 96 F.3d 1398 (Fed. Cir. 1996).....	5
<i>In re Katz Interactive Call Processing Patent Litigation</i> , 639 F.3d 1303 (2011).....	28
<i>Intamin Ltd. v. Magnetar Techs., Corp.</i> , 483 F.3d 1328 (Fed. Cir. 2007).....	19
<i>Intel Corp. v. Tela Innovations, Inc.</i> , No. 3:18-CV-02848-WHO, 2019 WL 5697922 (N.D. Cal. Nov. 4, 2019)	19
<i>Interval Licensing LLC v. AOL, Inc.</i> , 766 F.3d 1364 (Fed. Cir. 2014).....	26
<i>IPXL Holdings, L.L.C. v. Amazon.com, Inc.</i> , 430 F.3d 1377 (Fed. Cir. 2005).....	28
<i>Kaken Pharm. Co. v. Iancu</i> , 952 F.3d 1346 (Fed. Cir. 2020).....	10

<i>KEYnetik, Inc. v. Samsung Elecs. Co., Ltd.</i> , 837 F. App'x 786 (Fed. Cir. 2020)	10
<i>Keystone Bridge Co. v. Phoenix Iron Co.</i> , 95 U.S. 274 (1877).....	9
<i>Nautilus, Inc. v. Biosig Instruments, Inc.</i> , 572 U.S. 898 (2014).....	19, 21, 26
<i>Personalized Media Commc'ns, LLC v. Apple Inc.</i> , 952 F.3d 1336 (Fed. Cir. 2020).....	11, 14
<i>Saffran v. Johnson & Johnson</i> , 712 F.3d 549 (Fed. Cir. 2013).....	14
<i>Samsung Elecs. Am., Inc. v. Prisua Eng'g Corp.</i> , 948 F.3d 1342 (Fed. Cir. 2020).....	28, 29
<i>SCVNGR, Inc. v. DailyGobble, Inc.</i> , No. 6:15-CV-493-JRG-KNM, 2017 U.S. Dist. LEXIS 158494 (E.D. Tex. Sep. 26, 2017)	5

TABLE OF ABBREVIATIONS

Abbreviation	Term
Micron	Defendants Micron Technology, Inc., Micron Semiconductor Products, Inc., and Micron Technology Texas, LLC
Western Digital	Defendants Western Digital Corporation, Western Digital Technologies, Inc., and HGST, Inc.
Defendants	Defendants Micron and Western Digital
298 patent	U.S. Patent No. 8,891,298
385 patent	U.S. Patent No. 9,196,385
240 patent	U.S. Patent No. 9,997,240
300 patent	U.S. Patent No. 10,950,300

Throughout the brief, where a term is argued as “Defendants’ Position,” it is being proposed by both defendants, Micron and Western Digital. However, Micron takes no position on the terms being presented as “Western Digital’s Position,” and Western Digital takes no position on the terms being presented as “Micron’s Position.”

I. **INTRODUCTION**

There are good reasons behind this Court’s practice of construing claim terms well before trial. “This practice enables the parties to understand the key issues for trial more thoroughly and can guide the preparation of expert reports. In addition, claim construction can provide the basis for summary judgment, potentially simplifying or resolving the case. Early claim construction also enables the court and the parties to begin to prepare jury instructions.” Patent Case Management Judicial Guide (3d ed. 2016), at 2-13. Construing claims also makes them clearer for jurors.

Despite the many recognized advantages of early claim construction, Vervain, LLC (“Vervain”) refuses to meaningfully participate in the claim construction process. For most terms for which Micron and Western Digital seek a construction, Vervain proposes nothing more than a “plain and ordinary meaning.” Yet the parties’ discussion during the meet-and-confer over claim construction confirmed that there are real disputes over what these terms mean, including over what the so-called the “plain and ordinary meaning” is. Vervain’s approach does nothing to address or resolve those disputes, which is contrary to Federal Circuit precedent. Micron and Western Digital respectfully request that the Court adopt their approach and their proposed constructions.¹

II. **BACKGROUND**

The patents-in-suit relate to algorithms for improving the reliability of memory devices comprising both single-level cell (“SLC”) and multi-level cell (“MLC”) flash memory. SLC and MLC have known advantages and disadvantages relative to the other. All flash memory cells wear

¹ Exhibits cited herein are attached to the Declaration of J. Jason Lang, filed concurrently herewith. The patents-in-suit are attached as Ex. 1 (298 patent), Ex. 2 (385 patent), Ex. 3 (240 patent), and Ex. 4 (300 patent).

out; that is, they become defective after a certain number of write/erase cycles.² But SLC has greater endurance; it can sustain many more erase operations than MLC before becoming defective. SLC can also be written to more quickly than MLC. On the other hand, MLC can store more data per unit cost than SLC. 298 patent, 2:8–9;³ Declaration of Joseph C. McAlexander III in Support of Defendants’ Opening Claim Construction Brief (“McAlexander Decl.”, filed concurrently herewith) ¶¶ 21–23.

The patents-in-suit use logical-to-physical address mapping. The location in memory where data is stored is identified using a number known as an “address.” McAlexander Decl. ¶¶ 24; Ex. 5 (Microsoft Computer Dictionary) at 19 (definition of “address”). Addresses can be logical or physical. The “logical” address space is the set of addresses that the higher-level system, that is, the host, uses. Thus, for example, operating systems (*e.g.*, Microsoft Windows) or software (*e.g.*, Microsoft Word) would use logical addresses to identify locations within a memory to read or write. The “physical” address space is the set of addresses that identifies a specific physical cell or set of cells within the memory. Before each access to a flash memory occurs, the controller of the memory “translates” a logical address provided by the host into a physical address corresponding to specific flash memory cells. This translation is performed using a table called an “address map” that maintains a mapping between logical addresses and physical addresses. As the patents’ background section admits, logical-to-physical address mapping was well known in the art. 298 patent, 2:65–3:1; McAlexander Decl. ¶ 25.

The patents-in-suit purport to teach two techniques for improving the reliability of flash devices that have both MLC and SLC memory modules. 298 patent, 3:38–45. In the first

² A block of flash memory must be erased before it is written to again. Because of the relatively high voltages involved, this erase operation causes wear to the flash memory cell.

³ The other patents-in-suit share the same specification and, therefore, also support this proposition.

technique, the system performs a “data integrity test” on data stored ***in an MLC module*** and, if the test fails, the system transfers the data ***to an SLC module***. 298 patent, 5:34–40. In the second technique, the controller transfers blocks stored in MLC that receive frequent writes (*i.e.*, “hot blocks”) into SLC. 298 patent, 6:24–29. In both techniques, the transfer of data from an MLC memory module to an SLC memory module is performed by moving the data to the SLC memory module and updating the logical-to-physical address map to reflect the new physical address of the data in the SLC memory module. 298 patent, 5:50–54.

Underlying both techniques is a desire to combine the best aspects of SLC and MLC; the endurance of SLC and the cheap storage capacity of MLC. 298 patent, 1:25–32. Generally, data should be stored in the cheaper, higher-density storage medium—*i.e.*, MLC. But each disclosed technique serves to identify data worth storing in SLC. First, a block of data that has failed a data integrity test once might be more likely to fail it again, and so should be moved to the higher-endurance storage medium—*i.e.*, SLC. McAlexander Decl. ¶ 22. Second, a block of data that has been frequently written in the past is more likely to be written in the future, and so should be moved to the higher-speed and higher-endurance storage medium—*i.e.*, SLC. *Id.* Thus, the patents state that the disclosed “NAND flash storage system . . . provides long lifetime (endurance) storage at low cost.” 298 patent, 3:43–45.

III. ARGUMENT

A. “SLC non-volatile memory” (298 patent, cls. 1, 5; 385 patent cls. 1, 5; 240 patent, cls. 1, 6); “SLC nonvolatile memory” (300 patent, cls. 1, 3, 4, 7, 12)

Western Digital’s Proposed Construction	Vervain’s Proposed Construction
Non-volatile memory where each cell is capable of storing no more than one bit of information per cell	Plain and ordinary meaning, where the plain and ordinary meaning is “nonvolatile memory that stores one bit of information per cell”

B. “MLC non-volatile memory” (298 patent, cls. 1, 4; 385 patent, cls. 1, 4; 240 patent, cls. 1, 6); “MLC nonvolatile memory” (300 patent, cls. 1, 3, 4, 12)

Western Digital’s Proposed Construction	Vervain’s Proposed Construction
Non-volatile memory where each cell is capable of storing multiple bits of information per cell	Plain and ordinary meaning, where the plain and ordinary meaning is “nonvolatile memory that stores multiple bits of information per cell”

The parties’ constructions are similar: SLC stores only one bit of information per cell, while MLC may store one bit or multiple bits of information per cell. Two issues underpin the parties’ dispute. First, whether MLC memory and SLC memory can be properly defined by the number of bits each happens to store at a particular moment in time. It cannot. Second, whether MLC memory that, by happenstance, is storing only a single bit of information per cell at a given moment somehow is converted from MLC to SLC memory. It is not.

The terms “SLC non-volatile memory” and “MLC non-volatile memory” are used in the same claim numerous times to denote two different and separate things which have known advantages and disadvantages relative to each other. Sechen Decl., ¶¶ 31–33. As an initial matter, SLC and MLC are acronyms for “single level cell” and “multi-level cell” respectively. This alone confirms Western Digital’s construction which distinguishes the two types of memory by whether they are limited to a single bit per cell (SLC) or permit multiple bits per cell (MLC). Sechen Decl., ¶¶ 36–37. Moreover, contrary to Vervain’s construction, the claims make clear that SLC and MLC memory are different things and not interchangeable. For example, claim 1 of the 298 patent

claims “at least one MLC non-volatile memory module comprising a plurality of individually erasable blocks,” and then claims as a separate element, “at least one SLC non-volatile memory module comprising a plurality of individually erasable blocks.” Beyond claim 1, every independent claim in each of the asserted patents separately claims SLC memory and MLC memory as discrete and different things. *See*, 385 patent, cl. 1 and 6; 240 patent, cl. 1 and 6; 300 patent, cl. 1 and 12.

This usage in the claim language invokes two closely-related canons of claim construction that compel Western Digital’s construction and refute Vervain’s. First, “In the absence of any evidence to the contrary, we must presume that the use of these different terms in the claims connotes different meanings.” *CAE Screenplates, Inc. v. Henrich Fiedler GMBH & Co. KG*, 224 F.3d 1308, 1317 (Fed. Cir. 2000); *see also*, e.g., *SCVNGR, Inc. v. DailyGobble, Inc.*, No. 6:15-CV-493-JRG-KNM, 2017 U.S. Dist. LEXIS 158494, at *11 (E.D. Tex. Sep. 26, 2017) (finding that “distinct recitals of a ‘code’ and a ‘token’ in the same claim weigh in favor of finding that the ‘code’ and the ‘token’ must be distinct from one another”). Second, “[w]here a claim lists elements separately, ‘the clear implication of the claim language’ is that those elements are ‘distinct components’ of the patented invention.” *Becton, Dickinson & Co. v. Tyco Healthcare Grp., LP*, 616 F.3d 1249, 1254 (Fed. Cir. 2010); *see also*, e.g., *Bell Atl. Network Servs., Inc. v. Covad Commc’ns Group, Inc.*, 262 F.3d 1258, 1272–73 (Fed. Cir. 2001) (construing “mode” and “rate” to have different meanings because patent used those terms “to refer to two different and distinct concepts”); *Engel Indus., Inc. v. Lockformer Co.*, 96 F.3d 1398, 1404–05 (Fed. Cir. 1996) (holding “second portion” and “return portion” recited in the same claim could not “logically be one and the same”). Here, as noted above, the claims repeatedly cite to MLC and SLC memory by using different words to denote different and distinct components of the claimed inventions. Thus, there

can be no genuine question that they are different and discrete types of memory and an infringing product (or invalidating prior art reference) must have both types. Western Digital’s constructions recognize these important concepts by construing SLC and MLC memory by the critical distinction that separates them: the inability to store more than one bit per cell in SLC memory, and the ability to do so in MLC memory. Vervain’s construction, on the contrary, would violate these legal principles and effectively erase the distinction between MLC and SLC memory, allowing MLC memory to be re-categorized as SLC memory any time MLC memory happened to be storing only a single bit per cell. Just like a machinegun with a magazine does not convert to a single-shot rifle when one bullet is loaded in its magazine, neither does MLC convert to SLC memory when only a single bit of information happens to be present.

The specification too is dispositive and further confirms the correctness of Western Digital’s constructions and the error in Vervain’s when they describe, over and again, SLC and MLC as separate and different types of memory, and never describe MLC as somehow becoming SLC when only one bit is stored. For example, the specification describes “a plurality of MLC NAND flash memory modules . . . and a plurality of SLC NAND flash memory modules.” 298 patent, 6:40–44; *see also id.* at 1:25–32 (differentiating SLC and MLC); 1:38–43 (same); 1:64–67 (same); 3:15–30 (same); 3:33–36 (same); 4:10–15 (same); 4:53–55 (same); 5:13–15 (same); 5:24–27 (same). The specification then goes further in establishing that the difference between the two types of memory is exactly what Western Digital says it is; explaining that “[i]n SLC NAND flash technology, each cell ***can exist*** in one of two states [i.e., as a ‘0’ or a ‘1’], ***storing one bit of information per cell.***” *E.g.*, 298 patent, 1:64–65.⁴ In contrast, the specification describes MLC memory as “allow[ing] a single cell to ***store multiple bits.***” 298 patent, 3:15–16; *see also id.* at

⁴ Unless otherwise indicated, emphases in quotations from the specification or claims are added.

1:61–64 (“MLC NAND flash is a flash memory technology using ***multiple levels per cell to allow more bits to be stored*** using the same number of transistors.”).

The prosecution history also confirms Western Digital’s construction. During prosecution of the 300 patent, the applicant told the Patent Office that “‘MLC NAND flash is a flash memory technology using multiple levels per cell to allow more bits to be stored’ (relative to SLC).” Ex. 6 (300 Response to OA dated Mar. 23, 2020) at 6 (quoting the 300 patent application) (emphasis in original)). In other words, while MLC memory allows for the storage of multiple bits, it ***may*** store just one bit. Sechen Decl. ¶¶ 35–37.

Finally, Western Digital’s constructions are also supported by the extrinsic evidence. For example, a NAND flash textbook from around the time of the priority date of the patents defines SLC and MLC in terms of their capabilities:

- “MLC NAND is by far the lowest cost semiconductor memory with none of the memory technologies even close to being cost competitive. This is mainly due to the very small cell size combined with multi-level cell capability.” Ex. 7 (Micheloni) at 5.
- “Conventional SLC or single-level cell storage distinguishes between a ‘1’ and ‘0’ by having no charge or charge present on the floating gate of the Flash memory cell.” *Id.*

See also Sechen Decl. ¶¶ 33–34, 39.

Moreover, academic and industry articles contemporary to the asserted patents further confirm Western Digital’s construction, *e.g.*:

- “MLC chip technology is ***capable of*** storing two or more bits of data per memory cell, while SLC chip technology allows ***just one*** bit of data per memory cell.” Ex.

8 (Zhan et al.).

- “For SLC flash, ***only one bit could be stored*** to each cell at a time, and there will be two possible states for each cell – 0 or 1. As for MLC flash, ***two bits could be stored*** to each cell at a time and there will be four possible states for each cell – 00, 01, 10 or 11.” Ex. 9 (Chen et al.) at 3.

Sechen Decl. ¶¶ 38–39; *see also* Ex. 10 (Lee et al.) at 1–2 (2009 article by Lee et al. explaining that, “[i]n case of a single-level cell (SLC) flash memory, each cell has two states, and therefore ***only*** a single bit can be stored in that cell,” and consistently distinguishing MLC memory from SLC memory even when only one bit per cell is stored in MLC memory).

In sum, the evidence is overwhelming that SLC memory is capable of storing only a single bit of data per cell whereas MLC memory may store a single but or multiple bits per cell. Western Digital’s constructions should thus be adopted. That MLC memory may be programmed to store just one bit per cell does not transform it into SLC memory, and so Vervain’s construction should be rejected. Sechen Decl. ¶ 37.

C. “data integrity test” (298 patent, cl. 1; 385 patent, cl. 1; 240 patent, cls. 1, 6; 300 patent, cls. 1, 12)

Micron’s Proposed Construction	Vervain’s Proposed Construction
Plain and ordinary meaning, which is “a test conducted on data after it has been written to flash to ensure that the data was written correctly”	Plain and ordinary meaning.

The dispute between the parties is whether the “data integrity test” must be a test on data to ensure that the data was written correctly (Micron’s position), or whether a test to ensure that data was read from the correct address, *i.e.*, an address matching test, is ***also*** a data integrity test

(Vervain's position).⁵

Vervain's contention that an address matching test is a "data integrity test" is inconsistent with the plain meaning of the term. In its plain meaning, "data integrity" relates to whether data retains its intended value (was written correctly) or, instead, whether the data has been corrupted (was written incorrectly). The particular location in which data is stored, whether it is the intended location or another location, has nothing to do with the data's *integrity*. Vervain's attempt to stretch the claim beyond its plain meaning lacks merit. McAlexander Decl. ¶ 29.

Vervain's contention also is inconsistent with the patents' specification. The specification confirms that a "data integrity test" is a test "conducted on data . . . to ensure that the data was written correctly." The Summary of the Disclosure section explains that "*[a]fter each write* to (flash) memory, the controller conducts a data integrity check *to ensure that the data was written correctly.*" 298 patent, 4:4–6 (emphasis added). Micron's proposed construction of "data integrity test" as "a test conducted on data after it has been written to flash to ensure that the data was written correctly" tracks this language. The Detailed Description explains how the data integrity test is performed. It explains that "[d]ata to be written to NAND flash is maintained in DRAM 20. *After each write to an address* within a particular address range, the device controller 14 will—as time permits—perform a read on the address range *to ensure the integrity of the written data.* If a data integrity test fails, the address range is remapped from the MLC NAND flash memory module 26 to the next available address range in the SLC NAND flash memory module 28." 298 patent, 5:33–40. The plain meaning of "data integrity test" does not include tests to check whether data was read from the correct location, and nowhere does the specification suggest an intent to override

⁵ Vervain's contentions were made clear from its infringement contentions and the parties' meet and confer.

that meaning. McAlexander Decl. ¶ 31. “When the terms of a claim in a patent are clear and distinct (as they always should be), the patentee, in a suit brought upon the patent, is bound by it. He can claim nothing beyond it.” *Keystone Bridge Co. v. Phoenix Iron Co.*, 95 U.S. 274, 278 (1877). The term “data integrity test” as used in the patent is clear. Vervain cannot now go beyond the term’s plain meaning (and the specification) to include tests that ensure that data was read from the correct address.

Vervain’s interpretation is also unsupported by the purpose of the claimed invention. The specification discloses performing a data integrity test on data stored in an MLC module and, if the test fails, the system transfers the data to an SLC module. 298 patent, 5:34–40. The technique seeks to move data that has lost integrity (*i.e.*, failed a data integrity test) to the more reliable storage medium—that is, SLC—thereby reducing the chance of further data integrity test failures. McAlexander Decl. ¶¶ 32–34. But moving data to SLC if it turns out that the data was read from the wrong address—which, to Vervain, is a data integrity test failure—does not reduce the likelihood of reading data from the wrong address in the future. *Id.* Instead, a POSA would have understood that address matching test failures and data integrity test failures are fixed in different ways. McAlexander Decl. ¶ 35. Vervain’s artificial expansion of “data integrity test” “would eviscerate the stated purpose of the claimed invention.” *See KEYnetik, Inc. v. Samsung Elecs. Co., Ltd.*, 837 F. App’x 786, 792 (Fed. Cir. 2020) (citing *Kaken Pharm. Co. v. Iancu*, 952 F.3d 1346, 1352 (Fed. Cir. 2020)).

Because the intrinsic record is clear, the Court need not consult extrinsic evidence. Even so, the extrinsic evidence identified by Vervain does not support its interpretation. The Microsoft Dictionary defines “data integrity” as “[t]he accuracy of data and its conformity to its expected value, especially after being transmitted or processed.” Ex. 5 (Microsoft Computer Dictionary) at

143. The Hargrave's Dictionary defines "data integrity" as "[t]he condition that exists when data are unaltered after a process as compared to data before the process." Ex. 11 (Hargrave's Communications Dictionary) at 138. Vervain's contention—that a test to ensure that the data was read from the right place (*i.e.*, from the correct address) is a "data integrity test"—is supported by neither of these definitions. McAlexander Decl. ¶ 30.

D. "comparing the stored data to the retained data in the random access volatile memory" (300 patent, cls. 1 and 12)

Defendants' Proposed Construction	Vervain's Proposed Construction
comparing the data obtained by reading the nonvolatile memory space to the data retained as part of a Write access operation, wherein both sets of the data are in the same random access volatile memory	Plain and ordinary meaning

Defendants' construction of the "comparing" limitation is compelled by Vervain's argument regarding the scope of its claims during prosecution of the 300 patent. In order to convince the Patent Office to issue its claims over cited prior art, Vervain amended its claims to include the "comparing" limitation, and told the Patent Office that its claimed comparison requires comparing two specific sets of data—namely, the "retained data" to the data read from the nonvolatile memory just written to—that are both located in the same random access volatile memory. Those actions necessarily inform the meaning of the subject limitation, and disclaim any alternative system. *See, e.g., Personalized Media Commc'nns, LLC v. Apple Inc.*, 952 F.3d 1336, 1340 (Fed. Cir. 2020) ("[A]n applicant's amendment accompanied by explanatory remarks can define a claim term by demonstrating what the applicant meant by the amendment."). Vervain's prosecution history arguments—and Defendants' construction of the "comparing" limitation—are also consistent with the patent specification.

The claims in Vervain's initial application for the 300 patent did not include a data integrity

test at all. Ex. 12 (300 Application dated June 6, 2018) at 14–15. After a first rejection, Vervain amended its proposed claims to require a data integrity test without reciting how the data integrity test must be performed. Ex. 13 (300 OA dated Jan. 24, 2019) at 3–7; Ex. 14 (300 Response to OA dated June 26, 2019) at 2–5. The examiner again rejected the proposed claims as, among other things, obvious in view of several prior art references, including one to Oribe (U.S. Patent Application Publication No. 2009/0172267 A1). Ex. 15 (300 OA dated Dec. 13, 2019) at 4–12. With respect to the data integrity test limitation, the examiner found that Oribe taught using a data integrity test, “i.e., bit error greater than threshold,” to “mark a block for refresh.” *Id.* at 6, 12. Oribe further taught “that in a refresh the data is moved to a new block and the address management table is updated.” *Id.* The examiner explained that it would have been obvious to “use the data integrity test to ensure that the data is saved.” *Id.*

Vervain then amended its claims to require the “transfer of data” determined to have failed the data integrity test. Ex. 6 (300 Response to OA dated March 23, 2020) at 2, 5. Vervain argued that the operation in Oribe is “considerably different than that associated with the currently claimed process.” *Id.* at 12. First, Vervain argued that the data integrity test of the claims “requires a retention of the data **written to the Flash**,” an “**important distinction**” from Oribe. *Id.* at 11. Second, Vervain contrasted its data integrity test with Oribe’s “bit error test.” *Id.* at 12. In Oribe, according to Vervain, an error check code is created during a data write operation. *Id.* at 13. When that error check code is read, “a determination as to utilizing this error check code to correct the error can be made.” *Id.* at 13. Vervain asserted that Oribe’s bit error test “is **not** a compare operation with the data that was read but, rather, it is a determination that, if the integrity of the data has been compromised for some reason, it might be correctable depending upon how many bit errors are present.” *Id.* at 13. Accordingly, Vervain made clear that the data compared in its

claims is actual data written to Flash.

The examiner again rejected the claims as, among other things, obvious. Ex. 16 (300 OA dated April 28, 2020) at 5–11. In response, Vervain amended its proposed independent claims to how they read today, and argued that its amended “limitation requires that the nonvolatile memory element ‘retain’ the stored data and then, during a data integrity test, the stored data for a Write access is ‘compared’ to the stored data.” Ex. 17 (300 Response to OA dated Oct. 8, 2020) at 2, 5, 8.⁶

Vervain then provided a step-by-step description of the compare operation required by the claims:

1. The data that is being written to the Flash memory is “retained in nonvolatile memory.”
2. The data is “then written to memory,” i.e., Flash (nonvolatile) memory.
3. The data just written to Flash is “subsequently read from the memory.”
4. The read data is “compared to the retained data.”

Ex. 17 (300 Response to OA dated Oct. 8, 2020) at 9.

Vervain further explained that this compare operation cannot be carried out by a “conventional cache” because a “conventional cache” is “typically only designed to handle one block of data.” *Id.* In contrast, to facilitate the claimed compare operation, there must be at least two blocks of data in the volatile memory to store both the retained data and the read back data, i.e., the sets of data that are compared:

[I]f an entire block is being written to with new data, that entire block has to be retained in one portion of the memory and then written to the memory, followed by

⁶ Vervain mistakenly says here that the “nonvolatile” memory element must retain the data. Given that amended claims, in no uncertain terms, say that, in a Write access operation, the controller “retain[s] such stored data in the random access **volatile** memory,” Vervain must have meant to say that the **volatile** memory element retains the stored data. Sechen Decl. ¶ 46. Vervain makes this same mistake several more times in its response.

a subsequent Read of the memory, *requiring that block of data to be placed in an additional area of the nonvolatile memory, thus requiring at least two blocks of memory space in the nonvolatile memory.*

Id. Thus, the sets of data that are compared—i.e., the retained data (step 1 above) and data that is read from the Flash memory just written to (step 3 above)—must both be located in the same random access volatile memory. Sechen Decl. ¶ 48.

Vervain again contrasted this process with Oribe’s bit error test, which “merely analyze[s] the data and the error correction bits to determine the bit error rate”—not the “comparison operation” required by the claims. Ex. 17 (300 Response to OA dated Oct. 8, 2020) at 8–9. Moreover, Vervain argued that Oribe’s operation “does not in any way ‘retain’ data in any type of data cache for the purposes of a refresh operation.” *Id.* at 9.

Only after these arguments, the examiner allowed the claims. Ex. 18 (300 Jan. 27, 2021 Notice of Allowance).

A patentee “must be held to what he declares during the prosecution of his patent, because a contrary rule would undermine ‘[t]he public notice function of a patent.’” *Ajinomoto Co. v. ITC*, 932 F.3d 1342, 1351 (Fed. Cir. 2019) (citation omitted). Patentees “rarely submit affirmative disclaimers along the lines of ‘I hereby disclaim the following ...’ during prosecution and need not do so to meet the applicable standard.” *Saffran v. Johnson & Johnson*, 712 F.3d 549, 559 (Fed. Cir. 2013) (citation omitted). By its amendment and its clear and unequivocal statements that got these claims issued over repeated prior-art rejections, Vervain disclaimed any method of performing the comparison operation that (1) analyzes error bits, as opposed to the actual data to be written to Flash and (2) does not compare two sets of data—namely, the retained data to the data read from the Flash memory just written to—located in the same random access volatile memory. *See, e.g., Biogen Idec, Inc. v. GlaxoSmithKline LLC*, 713 F.3d 1090, 1094–97 (Fed. Cir.

2013); *Personalized Media*, 952 F.3d at 1346. Defendants' construction is consistent with that disclaimer. *See also* Sechen Decl. ¶¶ 43, 48.

The specification is consistent with the construction compelled by the prosecution history. For example, the specification describes performing the comparison only on the actual data to be written to Flash, and not on error bits. *See* 300 patent, 6:17–52; Sechen Decl. ¶¶ 47–48. The specification also makes clear that data to be written to Flash is retained only in the volatile memory. *See, e.g.*, 300 patent, 5:60–61, 6:18–19. And thus Vervain's prosecution statements that its comparison operation requires more cache space than a “conventional cache” only make sense if Vervain was referring to needing to store both the retained data (which is always in volatile memory) and the read back data in the volatile memory.

In sum, in order to get its 300 patent to issue, Vervain told the Patent Office that a bit error test was insufficient as a comparison, but rather that stored data (data obtained by reading the nonvolatile memory space) is compared to retained data (data retained as part of a Write access operation), wherein both sets of the data are in the same random access volatile memory. The Court should not allow Vervain to write those critical distinctions out of its claims now, and so should adopt Defendants' construction.

E. “to achieve enhanced endurance” (300 patent, cl. 1, 12)

Micron's Proposed Construction	Vervain's Proposed Construction
Plain and ordinary meaning, which is “to achieve endurance (i.e., lifetime) superior to that of the MLC nonvolatile memory element”	Plain and ordinary meaning.

Claims 1 and 12 of the 300 patent recite:

the controller performing a data integrity test on stored data in the MLC nonvolatile memory element . . .

wherein a failure of the data integrity test . . . results in . . . transfer of data corresponding to the stored data to those remapped physical addresses from those determined to have failed the data integrity test ***to achieve enhanced endurance.***

The parties' dispute turns on the circumstances in which a transfer of data "achieve[s] enhanced endurance." Micron contends that a transfer of data achieves enhanced endurance only when it is transferred from the MLC nonvolatile memory element (where, under the claim language, it was originally stored) to a storage medium with endurance (*i.e.*, lifetime) superior to that of the MLC nonvolatile memory element (*e.g.*, an SLC nonvolatile memory element). Vervain contends that a transfer of data can achieve enhanced endurance even when it remains in the MLC nonvolatile memory element, such as when it is moved from a location in MLC to another, less-worn location in MLC.

Micron's proposed construction follows from the plain language of the claims. In claims 1 and 12 of the 300 patent, "the controller perform[s] a data integrity test on stored data in the MLC nonvolatile memory element." "[A] failure of the data integrity test . . . results in . . . transfer of data corresponding to the stored data . . . ***to achieve enhanced endurance.***" Because the data was originally stored in the MLC nonvolatile memory element, to achieve ***enhanced*** endurance is to achieve endurance superior to that of the MLC nonvolatile memory element. McAlexander Decl. ¶ 39.

Vervain's position—that a transfer of data from one MLC location to another, potentially less worn location within the MLC nonvolatile memory element "achieve[s] enhanced endurance"—lacks merit. The specification repeatedly and consistently uses the word "endurance" to contrast MLC, which is lower endurance, and SLC, which is higher endurance:

- "This application relates . . . to a system and method of increasing the reliability and lifetime of a NAND flash storage system . . . through the use of a combination of single-

level cell (SLC) and multi-level cell (MLC) NAND flash storage. . . . The SLC non-volatile memory can be flash . . . or any other solid-state non-volatile memory *as long as it has endurance that is superior to that of MLC flash* . . . or rotating storage media (e.g., HDDs).” 298 patent, 1:25–43; *see also id.* at 4:10–15 (Summary of the Disclosure section stating the same).

- “MLC NAND flash enjoys greater density than SLC NAND flash, at the cost of a *decrease in . . . lifetime (endurance)*.” 298 patent, 3:19–21.
- “[T]wo separate banks of NAND flash are maintained by a controller. One bank contains economical MLC NAND flash, while a second bank contains *high endurance SLC NAND flash*.” 298 patent, 4:52–55.
- “As *the SLC NAND flash is used to boost the lifetime (endurance) of the storage system*, it can be considerably lesser in amount than the MLC NAND flash.” 298 patent, 4:58–61.

McAlexander Decl. ¶ 41–42. The specification always describes “endurance” as a property of a type of storage (e.g., MLC, SLC, hard-disk drives, or “flash alternatives”). 298 patent, 2:6–9, 3:22–27. Nowhere does the specification suggest that lower-wear MLC blocks are higher “endurance” MLC blocks. This despite the background section’s discussion of static and dynamic wear leveling, both of which move data to lower-wear blocks. What’s more, the specification makes clear that the *reason* SLC is higher endurance than MLC is that SLC has greater lifetime—that is, SLC can withstand more write/erase cycles than MLC.

- The specification notes that relative to a flash memory, “flash alternatives . . . possess superior endurance (*1,000,000 write-erase cycles* compared to typical *100,000 cycles* for SLC flash and less than *10,000 cycles* for MLC flash).” 298 patent, 2:6–9.

- In comparing flash to magnetic media, the specification notes that “even SLC NAND flash is considerably lower lifetime (endurance) than rotating magnetic media (e.g., HDDs), being able to withstand only between ***50,000 and 100,000 writes***, and MLC NAND flash has a much lower lifetime (endurance) than SLC NAND flash, being able to withstand only between ***3,000 and 10,000 writes***.” 298 patent, 3:22–27.

McAlexander Decl. ¶ 40. The specification’s use of “endurance” to mean the write-erase lifetime of a type of storage reflects the ordinary and customary meaning of “endurance” to a POSA:

- Multiple JEDEC standards documents define the “endurance (of a reprogrammable read-only memory)” as “[t]he ability of a ***reprogrammable read-only memory*** to withstand ***data rewrites*** and still comply with its specifications.” Ex. 19 (JEDEC Standard 100B.01) at 4; Ex. 20 (JEDEC Standard 22A117B) at 2.
- A textbook on NAND flash memory states that “[t]he endurance of a ***memory module*** is defined as minimum ***number of Program/Erase cycles*** that the module can withstand before leading to a failure.” Micheloni at 91.

McAlexander Decl. ¶ 38. In sum, a POSA would have understood the word “endurance” to refer to the number of write-erase cycles that a type of storage can sustain. A transfer of data from one location to another within an MLC nonvolatile memory element cannot “achieve enhanced endurance” because the MLC nonvolatile memory element cannot have a greater write-erase-cycle lifetime than itself. McAlexander Decl. ¶ 43.

Vervain’s interpretation of “enhanced endurance” also conflicts with the written description support for and the purpose of the claimed invention. The specification discloses performing a data integrity test on data stored in an MLC module and, if the test fails, the system transfers the data to an SLC module. 298 patent, 5:34–40. The purpose of the technique is to

move data that has lost integrity (*i.e.*, failed a data integrity test) to the more reliable storage medium—that is, SLC—thereby reducing the chance of further data integrity test failures. McAlester Decl. ¶ 40. Indeed, the specification explains that “the SLC NAND flash is used to boost lifetime (endurance) of the storage system.” 298 patent, 4:58–59. Because a transfer of data to another address in MLC would not “boost [the] lifetime (endurance) of the storage system,” such a transfer would be inconsistent with the purpose of the claimed invention. *See Kaken*, 952 F.3d at 1352.

**F. “the list of logical address ranges having a minimum quanta of addresses”
(298 patent, cl. 1; 385 patent, cl. 1; 240 patent, cl. 1, 6)**

Western Digital’s Proposed Construction	Vervain’s Proposed Construction
Indefinite	Plain and ordinary meaning

The term “the list of logical address ranges having a minimum quanta of addresses” is indefinite because a person of ordinary skill in the art would not be reasonably apprised what that minimum is. The term “having a minimum quanta of addresses” is a term of degree, which must have an objective boundary to be definite. *See Berkheimer v. HP Inc.*, 881 F.3d 1360, 1363 (Fed. Cir. 2018) (finding “minimal redundancy” indefinite because the “claim language is not reasonably clear as to what level of redundancy in the archive is acceptable”); *Intel Corp. v. Tela Innovations, Inc.*, No. 3:18-CV-02848-WHO, 2019 WL 5697922, at *12 (N.D. Cal. Nov. 4, 2019) (emphasis omitted) (finding “minimum size” indefinite because intrinsic evidence “provides no objective way to determine what ‘minimum size’ means”). Here, there is no objective way to determine how many addresses are “minimum quanta.” The term thus fails to “inform those skilled in the art about the scope of the invention with reasonable certainty.” *Nautilus, Inc. v. Biosig Instruments, Inc.*, 572 U.S. 898, 910 (2014).

The claims provide no help. Indeed, the dependent claims further demonstrate why the

independent claims provide no objective boundary. While the independent claims require *some* “minimum quanta of addresses” (*see* 298 and 385 patents, claim 1; 240 patent, claims 1 and 6), dependent claim 2 of the 298 and 385 patents, for example, further requires that the “minimum quanta of addresses is equal to one block.” Claim 1 must be broader than claim 2. *See, e.g., Alcon Rsch., Ltd. v. Apotex Inc.*, 687 F.3d 1362, 1367 (Fed. Cir. 2012) (“It is axiomatic that a dependent claim cannot be broader than the claim from which it depends.”); *Intamin Ltd. v. Magnetar Techs., Corp.*, 483 F.3d 1328, 1335 (Fed. Cir. 2007) (“An independent claim impliedly embraces more subject matter than its narrower dependent claim.”). Therefore, the “minimum quanta of addresses” permitted by claim 1 must be *smaller than* one block, but how much smaller? The claims provide no answer. *See also* 298 and 385 patents, claim 3 (“The system of claim 1, wherein the minimum quanta of addresses is equal to one page.”).

The specification is also of no help. The specification references a “minimum” quanta of addresses only once. But, like the claims, the specification provides no information on what the minimum is, saying that the address ranges “will assume *some* minimum quantum.” *E.g.,* 298 patent, 5:27–31 (emphasis added). The specification goes on to provide the same examples required by the further limitations of dependent claims 2 and 3: “such as, for example, one block, although a smaller size, such as one page could be used, if the NAND flash has the capability of erasing the smaller size quantum.” *Id.* This does not inform a skilled artisan of the meaning of “minimum quanta of addresses” with reasonable certainty. *See* Sechen Decl. ¶¶ 53–55.

When faced with Western Digital’s position that this term is indefinite, Vervain did nothing to attempt to define it in its claim construction proposals, other than to state that this term should get its “plain and ordinary meaning.” But this invites Vervain to argue *anything* can be a “minimum quantum of addresses” for infringement purposes. If the term is not indefinite, it must

mean something.

- G. “wherein the controller is further adapted to determine which of the blocks of the plurality of the blocks in the MLC and SLC non-volatile memory modules are accessed most frequently and wherein the controller segregates those blocks that receive frequent writes into the at least one SLC non-volatile memory module and those blocks that receive infrequent writes into the at least one MLC nonvolatile module, and maintain a count value of the blocks in the MLC non-volatile memory module determined to have received frequent writes and that are accessed most frequently on a periodic basis when the count value is a predetermined count value, transfer the contents of the counted blocks in the MLC non-volatile memory module determined to have received frequent writes after reaching the predetermined count value to the SLC non-volatile memory module and which determined blocks in the SLC are determined in accordance with the next equivalent range of physical addresses determined by the controller” (240 patent, cl. 1)

Defendants’ Proposed Construction	Vervain’s Proposed Construction
Indefinite. Although it is possible for a person of ordinary skill (“POSA”) to identify some embodiments that fall within the scope of the claim, a POSA would not be reasonably certain of the full scope of the claim.	Plain and ordinary meaning.

The above limitation is present in claim 1 of the 240 patent. As Defendants explain below, multiple phrases within this limitation render the bounds of the limitation unclear. More specifically, the limitation creates three zones in the mind of a POSA: a zone in which systems clearly *are* practicing the claim; a zone in which systems clearly *are not* practicing the claim; and a “zone of uncertainty.” *Nautilus*, 572 U.S. at 909. For systems that fall into the zone of uncertainty, a POSA would not be reasonably certain as to whether the system is—or is not—practicing the claim. Each of three issues, standing alone, renders the full scope of claim indefinite.

First, the phrase “on a periodic basis” creates a zone of uncertainty as to when a system is sufficiently periodic. The limitation recites:

wherein the controller is further adapted to . . . maintain a count value of the blocks . . . determined to have received frequent writes and that are accessed most frequently ***on a periodic basis*** when the count value is a predetermined count value, transfer the contents of the counted blocks

A POSA would have understood that a controller practices this limitation when it performs the recited “transfer” on a periodic basis. McAlexander Decl. ¶ 49. But the limitation leaves uncertain whether other forms of periodicity also practice the limitation: (1) Does a controller practice this limitation (or not) if it determines “when the count value is a predetermined count value” on a periodic basis? (2) Does a controller practice this limitation (or not) when the controller determines the “blocks . . . that are accessed most frequently” on a periodic basis? (3) Is this limitation practiced (or not) when the “count value of the blocks . . . accessed most frequently” is a count of blocks accessed “most frequently on a periodic basis”? Or, (4) if multiple of these actions are within the scope of the limitation, which ones? In light of the myriad ways in which “on a periodic basis” can be read to modify different actions recited in the limitation, a POSA would not have been reasonably certain as to which actions the claim requires—and does not require—to be performed “on a periodic basis.” *Id.*

Second, the limitation includes two clauses that interact in a manner that renders the full scope of the claim unclear. The limitation begins with a first clause (the “segregation clause”), which states:

wherein the controller is further adapted to determine which of the blocks of the plurality of the blocks in the MLC and SLC non-volatile memory modules are accessed most frequently and wherein the controller segregates those blocks that receive frequent writes into the at least one SLC non-volatile memory module and those blocks that receive infrequent writes into the at least one MLC nonvolatile module.

Viewing the segregation clause in isolation, a POSA would have understood the clause to be practiced when the controller “segregates” frequently written blocks from MLC to SLC and

“segregates” infrequently written blocks from SLC to MLC. McAlexander Decl. ¶ 51

The limitation continues, however, with a second clause (the “transfer clause”), which states:

wherein the controller is further adapted to . . . maintain a count value of the blocks in the MLC non-volatile memory module determined to have received frequent writes and that are accessed most frequently on a periodic basis when the count value is a predetermined count value, transfer the contents of the counted blocks in the MLC non-volatile memory module determined to have received frequent writes after reaching the predetermined count value to the SLC non-volatile memory module.

Viewing the transfer clause in isolation, a POSA might have understood the clause to require transferring frequently written blocks from MLC to SLC. McAlexander Decl. ¶ 52. And, viewing both clauses together, a POSA would have understood that a system that separately practices both clauses is within the scope of the limitation. *Id.*

But the claim language does not make reasonably clear whether the *same* action can satisfy both the segregation clause and the transfer clause. On its face, the patentee’s decision to separately claim the “segregat[ing]” step and the “transfer” step suggest to a POSA that the steps require distinct actions. *Id.* But a POSA could also reasonably read the two clauses to be satisfied simultaneously. This is because a POSA could reasonably read the antecedent basis for the transfer clause’s phrase “blocks . . . **determined** to have received frequent writes” to be in the segregation clause, which recites “**determining** which of the blocks of the plurality of the blocks in the MLC and SLC non-volatile memory modules are accessed most frequently.” This renders the transfer clause indefinite because it is uncertain whether the transfer clause’s “blocks” refer to the “blocks” in the segregation clause or refer to a different set of blocks. In sum, a POSA would not have been reasonably certain of whether a system that satisfies the segregation clause inherently satisfies the transfer clause, or whether the clauses require discrete

operations on different sets of blocks, rendering the full scope of the limitation unclear. *Id.*

Third, the full scope of the phrase “a count value of the blocks . . . determined to have received frequent writes and that are accessed most frequently” would have been unclear to a POSA. Consider a hypothetical system (an “individual write counts” system) in which the controller maintains individual write counts for each of the plurality of blocks in the MLC non-volatile memory. In this system, blocks with the highest write counts are determined to have received frequent writes. This system would seem to make intuitive sense because, to determine which blocks “have received frequent writes,” a controller would need maintain individual counts of each block. In another hypothetical system (a “collective write count” system), the controller maintains a single, collective “count value” representing the total number of writes to **multiple** blocks in the MLC non-volatile memory. This system is consistent with the claim’s singular “count value.” A POSA would have understood an individual write counts system to be within the scope of the claims, and would also have understood a collective write count system to be within the scope of the claim. McAlexander Decl. ¶ 53.

But a POSA could reasonably read the claims to encompass other systems and would not be reasonably certain of whether the claims included—or excluded—each of these systems. Consider the phrase “a count value **of the blocks**.**”** The plain meaning of “of the blocks” would suggest that the “count value” is a count of a number of blocks (e.g., 10 blocks or 20 blocks). Combined with the alternative readings of “count” as individual and collective, a POSA could potentially envision both an “individual block count” system and a “collective block count” system as being within the scope of the limitation. Both systems would be different from the “individual write counts” and the “collective write count” systems discussed above, since those systems count **writes**, not blocks. Though arguably more consistent with the plain reading of the phrase “count

value of the blocks,” a *block* count would be inconsistent with the hot-blocks limitations of the other patents-in-suit, which instead recite an *access* (*i.e.*, *write*)⁷ count in claiming the same technique, and the specification does not hint at a block count. 298 patent, cl. 1; 385 patent, cl. 1; 298 patent, 6:24–35; McAlexander Decl. ¶ 54.

A POSA would have understood the claims to include individual write counts systems and collective write count systems to meet the limitation at issue. But the ambiguities outlined above prevent a POSA from being reasonably certain as to whether the limitation would be met (or not met) by: (1) an individual block counts system; and/or (2) a collective block count system. McAlexander Decl. ¶ 55.

H. “wherein the controller is further adapted to maintain a count value of those blocks that are accessed most frequently and, on a periodic basis when the count value is a predetermined count value, transfer the contents of those counted blocks into the SLC non-volatile memory module, wherein the counted blocks transferred to after reaching the predetermined count value are determined in accordance with the next equivalent range of physical addresses determined by the controller” (240 patent, cl. 6)

Defendants’ Proposed Construction	Vervain’s Proposed Construction
Indefinite. Although it is possible for a person of ordinary skill (“POSA”) to identify some embodiments that fall within the scope of the claim, a POSA would not be reasonably certain of the full scope of the claim.	Plain and ordinary meaning.

As with the previous limitation, multiple phrases within this limitation render the full scope of the limitation unclear. And, as with the previous limitation, this limitation creates a zone in which systems clearly are practicing the claim; a zone in which systems clearly are not practicing the claim; and a zone of uncertainty in which a POSA would not be reasonably certain as to

⁷ The patents use “write” and “access” interchangeably. See 298 patent, 6:30–32.

whether a system is—or is not—practicing the claim. Two issues create this zone of uncertainty.

First, the full scope of the phrase “a count value of the blocks . . . determined to have received frequent writes and that are accessed most frequently” creates a zone of uncertainty for the reasons explained above for claim 1 of the 240 patent. *See Section III.G, supra*; McAlexander Decl. ¶ 58.

Second, the bounds of the phrase “wherein the counted blocks transferred to after reaching the predetermined count value” is not sufficiently defined. In view of the specification, a POSA would have understood that transferring the counted blocks to SLC practices this limitation. McAlexander Decl. ¶ 59; 298 patent, 5:33–40. This understanding is supported by another part of the limitation, which states the “**contents** of those counted blocks” are “transfer[red] . . . into the SLC non-volatile memory module.” But that phrase describes the “contents” of the counted blocks, rather than the “counted blocks” themselves. The phrase “the counted blocks transferred to” does not unambiguously specify where (and where **not**) the blocks are “transferred to.” As to whether the limitation can be met when the counted blocks are transferred to somewhere other than SLC—*e.g.*, to other locations within the MLC non-volatile memory module—a POSA would not be reasonably certain. *Id.*

Each of these ambiguities, standing alone, prevents a POSA from having a reasonably certain understanding of the limitation’s full scope.

I. “wherein the mapping is performed as necessitated by the system to maximize lifetime” (300 patent, cls. 1, 12)

Western Digital’s Proposed Construction	Vervain’s Proposed Construction
Indefinite	Plain and ordinary meaning

The term “as necessitated by the system to maximize lifetime” is indefinite because it fails to inform a POSA about the scope of the alleged invention with reasonable certainty. *Nautilus*,

Inc. 572 U.S. at 910.

“[A]s necessitated by the system to maximize lifetime” is a term of degree, which must have an objective boundary to be definite. *See Berkheimer v. HP Inc.*, 881 F.3d at 1363 (term “wherein the archive exhibits minimal redundancy” “is not reasonably clear as to what level of redundancy in the archive is acceptable”). Although terms of degree are not inherently indefinite, terms of degree must still “provid[e] enough certainty to one of skill in the art when read in the context of the invention” and identify “some standard for measuring the scope of the phrase.” *Interval Licensing LLC v. AOL, Inc.*, 766 F.3d 1364, 1370–71 (Fed. Cir. 2014) (citation omitted). Here, the intrinsic evidence fails to provide a person skilled in the art with any objective way to determine how lifetime is maximized, how the lifetime must be increased to be “maximized,” or when such mapping to maximize lifetime is “necessitated.” Sechen Decl. ¶¶ 58–59, 61–62.

The claims provide no information useful to discern the scope of the term. Furthermore, the term “maximize lifetime” or any variation of “maximize” does not appear in the specification. The specification states that embodiments “relate to a system and method of *increasing* the reliability and *lifetime* of a NAND flash storage system, module, or chip through the use of a combination of multi-level cell (MLC) and single-level cell (SLC) NAND flash storage.” 300 patent, 7:16–20 (emphasis added). Additionally, the specification teaches “remap[ping] from MLC NAND Flash to SLC NAND Flash,” the “SLC NAND flash [being] used to *boost the lifetime* (endurance) of the storage system.” 300 patent, 5:15–20. But this is a far cry from “*maximizing* lifetime,” which is undefined. Nor does the specification explain when “mapping” would be “necessitated by the system to maximize lifetime” or how such mapping would “maximize lifetime.” Sechen Decl. ¶¶ 61–62. As mentioned, the specification discusses “remapping” (i.e., a separate claim limitation) to “boost the lifetime,” *see* 300 patent at 7:16–20, but does not say

anything about “mapping” to increase lifetime.

Finally, nothing in the intrinsic evidence gives clarity to the phrase “as necessitated by the system.” The specification does not disclose how or when the system would determine that mapping was necessary to maximize lifetime. Absent any disclosure in the specification to explain the condition(s) that the system would deem to necessitate “mapping … to maximize lifetime,” there is no standard upon which a skilled artisan could discern when the mapping would be “performed as necessitated by the system,” or when such mapping is necessary. Without that, the term “as necessitated by the system” merely discloses arbitrary mapping on an undefined as-needed basis.

J. “wherein a failure of the data integrity test performed by the controller results in a remapping of the address space to a different physical range of addresses and transfer of data corresponding to the stored data to those remapped physical addresses” (300 patent, cls. 1, 12)

Micron’s Proposed Construction	Vervain’s Proposed Construction
Indefinite under <i>IPXL Holdings, L.L.C. v. Amazon.com, Inc.</i> , 430 F.3d 1377 (Fed. Cir. 2005).	Plain and ordinary meaning.

As this Court is well aware, a system claim that recites method steps is indefinite. *See IPXL Holdings, L.L.C. v. Amazon.com, Inc.*, 430 F.3d 1377, 1384 (Fed. Cir. 2005). “The rationale of *IPXL* is that the claim conflates elements of both an apparatus and a method, rendering the claim indefinite for purposes of determining **when** infringement occurs.” *Samsung Elecs. Am., Inc. v. Prisua Eng’g Corp.*, 948 F.3d 1342, 1355 (Fed. Cir. 2020) (emphasis added). For example, in *IPXL*, the Court held the claims indefinite because “it is unclear whether infringement [of the claim] occurs **when** one creates a system [that allows a user to use the system] or whether infringement occurs **when** the user actually uses [the system].” 430 F.3d at 1384 (emphasis added). Similarly, in *Katz*, the Court held the claims at issue indefinite when they “create[d] confusion as to **when**

direct infringement occurs because they are directed both to systems and to actions performed by [humans].” *In re Katz Interactive Call Processing Patent Litigation*, 639 F.3d 1303, 1318 (2011) (emphasis added). Both cases stand for the proposition that a lack of reasonable certainty as to the *timing* of infringement renders a claim indefinite.

Claims 1 and 12 of the 300 patent both claim “[a] system for storing data.” Both claims recite: “wherein a failure of the data integrity test performed by the controller results in”: (1) “a **remapping** of the address space to a different physical range of addresses,” (the “remapping limitation”); and (2) “**transfer** of data corresponding to the stored data to those remapped physical addresses” (the “transfer limitation”). The parties disagree as to whether the term is indefinite.

The plain language of the limitations places no restriction on who or what performs the “remapping” or the “transfer.” McAlexander Decl. ¶ 64. Consider a hypothetical system, say a solid-state drive (“SSD”), in which the “flash translation layer” recited by the claim performs the “remapping” limitation. Does the SSD practice the limitation? Certainly. What if the “controller” recited by the claim performed the “remapping”? This would also practice the limitation. What if hardware *outside* the SSD—such as a host computer into which the SSD is plugged—performed the “remapping” limitation? By the limitation’s plain language, yes, this would practice the limitation. What if the “remapping” were performed manually by a human user? Again, the limitation as written would not exclude such an embodiment. The same goes for the “transfer” limitation; by its plain language, the limitation would include a transfer by any entity, human or machine, whether or not part of the claimed system. *Id.*

The open-ended language of the “remapping” and “transfer” limitations causes confusion as to *when*, as a matter of timing, the limitation is infringed. *Id.* ¶ 65. For instance, a system with an FTL or a controller configured to perform the remapping and transfer might infringe when it is

created. But if a host computer or human performs either the remapping or the transfer, or both, the system may not infringe until it is used. Because it is unclear “*when* infringement occurs,”—that is, at the time the system is manufactured or when it is used—the claims are indefinite. *See Samsung*, 948 F.3d at 1355.

Vervain’s position—that the “remapping” and “transfer” limitations merely establish the capability of the system—is belied by the plain language of the claim. The patentee knew full well how to claim capability. Both the 298 and 385 patents recite claims in which “a controller . . . *is adapted to*” perform various actions. Even the very claims that include the “remapping” and “transfer” limitations—claims 1 and 12 of the 300 patent—recite elsewhere that “the controller [is] *operable to* store data in the MLC nonvolatile memory element.” The patentee’s deliberate choice to use a passive, open-ended construction for the “remapping” and “transfer” limitations—instead of the “*operable to*” construction used elsewhere—must be given effect. McAlexander Decl. ¶ 66. “[C]ourts may not redraft claims, whether to make them operable or to sustain their validity.” *Chef Am., Inc. v. Lamb-Weston, Inc.*, 358 F.3d 1371, 1374 (Fed. Cir. 2004).

IV. CONCLUSION

For the reasons set forth above, the Court should adopt Defendants’ proposed constructions.

Respectfully submitted,

/s/ Douglas E. Lumish

Douglas E. Lumish (*pro hac vice*)
Richard G. Frenkel (*pro hac vice*)
Patricia Young (*pro hac vice*)
Miriam G. Cabello (*pro hac vice*)
LATHAM & WATKINS, LLP
140 Scott Drive
Menlo Park, CA 94025
Tel: (650) 328-4600
Fax: (650) 463-2600
rick.frenkel@lw.com
patricia.young@lw.com
doug.lumish@lw.com

Stephen A. Maniscalco (*pro hac vice*)
LATHAM & WATKINS, LLP
1271 Avenue of the Americas
New York, NY 10020
Tel: (212) 906-1200
Fax: (212) 751-4864
stephen.maniscalco@lw.com

Michael E. Jones
State Bar No. 10929400
mikejones@potterminton.com
Patrick C. Clutter
State Bar No. 24036374
patrickclutter@potterminton.com
POTTER MINTON, P.C.
110 N. College Ave., Suite 500
Tyler, Texas 75702
Tel: (903) 597-8311
Fax: (903) 593-0846

Attorneys for Defendants Western Digital Corporation, Western Digital Technologies, Inc., and HGST, Inc.

/s/ Jared Bobrow

Jared Bobrow (admitted *pro hac vice*)
CA Bar No. 133712
Jason Lang (admitted *pro hac vice*)
CA Bar No. 255642
ORRICK, HERRINGTON & SUTCLIFFE LLP
1000 Marsh Road
Menlo Park, CA 94025
Telephone: 650.614.7400
Facsimile: 650.614.7401
jbobrow@orrick.com
jlang@orrick.com

Claudia Wilson Frost

State Bar No. 21671300
ORRICK, HERRINGTON & SUTCLIFFE LLP
609 Main, 40th Floor
Houston, TX 77002
Telephone: 713.658.6400
Facsimile: 713.658.6401
cfrost@orrick.com

Margaret Abernathy (admitted *pro hac vice*)
CA Bar No. 300273
ORRICK, HERRINGTON & SUTCLIFFE LLP
777 South Figueroa Street, Suite 3200
Los Angeles, CA 90017
Telephone: 213.629.2020
mabernathy@orrick.com

Parth Sagdeo (admitted *pro hac vice*)
MA Bar No. 707696
ORRICK, HERRINGTON & SUTCLIFFE LLP
222 Berkeley St., Suite 2000
Boston, MA 02116
Telephone: 617.880.1800
Facsimile: 617.880.1801
psagdeo@orrick.com

Attorneys for Defendants Micron Technology, Inc., Micron Semiconductor Products, Inc., And Micron Technology Texas, LLC

CERTIFICATE OF SERVICE

The undersigned certifies that on November 12, 2021, all counsel of record who are deemed to have consented to electronic service are being served with a copy of this document through the Court's CM/ECF system. Any other counsel of record will be served by a facsimile transmission or first class mail.

/s/ Jared Bobrow
Jared Bobrow